The FAME Problem Domain For Distributed Planning

Markus Dale and **Marie desJardins**
{*dalem1,mariedj*}@*cs.umbc.edu*
Department of Computer Science
University of Maryland Baltimore County
1000 Hilltop Circle, Baltimore, MD 21250

Abstract

We introduce the Fortunate Artists Media Empire (FAME) problem domain to study approaches for solving key challenges in distributed planning such as dynamic environments, dynamic goals, uncertainty of actions and scalability. FAME is an artificial problem domain that allows us to focus on key domain-independent research issues. The FAME problem domain is focused on broadcast quality monitoring tasks for radio and TV stations. Besides the broadcast stations, the FAME infrastructure consists of a quality assurance (QA) monitoring headquarters, QA regional centers, and mobile, truck-mounted sensors that need to be employed effectively. An implementation of the FAME domain could serve as a testbed for distributed planning research.

Introduction

Controlling large-scale, complex distributed systems in dynamic, uncertain environments is a hard task for human operators. For example, creating a logistics plan for a major military operation involves many agencies; without automated tools, a large operation can take weeks to coordinate and plan. To apply automation to such tasks, we are interested in creating decision support systems that can help human operators. Automated planning—the ability to deduce sequences of actions that lead to desired goal states—is a key technology for such systems.

For dynamic, large-scale distributed systems in domains that exhibit a large amount of uncertainty, a centralized planning solution is often not desirable or even feasible. Privacy or security concerns may prevent sharing of sensitive data, but the participating parties can still coordinate a mutually acceptable plan in a distributed fashion by exchanging limited information. Likewise, large amounts of rapidly changing data in a network-constrained environment may preclude centralization of data. A distributed solution may also be desirable to allow for continuity of operations and survivability. Examples of such dynamic, large-scale distributed systems include parcel delivery services, disaster response logistics, and Navy carrier strike groups (CSGs). We explore elements of the CSG problem in further detail as an example of some of the key problems in distributed planning.

Figure 1: Abraham Lincoln Carrier Strike Group, Rim of the Pacific Exercise (RIMPAC) 2000 (Wikipedia 2007).

Distributed Planning for Carrier Strike Groups

A CSG typically consists of an aircraft carrier, a carrier air wing, a destroyer squadron, guided missile cruisers, submarines, and supply ships (U.S. Navy Office of Information 2004), as displayed in Figure 1.

CSGs are often assigned a prioritized list of mission objectives, which might not all be feasible, given the group's resource limitations. The CSGs also have to defend themselves against possible air, surface, and subsurface threats from an ever adjusting enemy. Weather, malfunctions, and human errors provide further uncertainty in this domain. Most of the assets in a CSG are multi-mission-capable. For example, a guided missile destroyer could be used in an anti-air, anti-surface, or anti-submarine warfare role. Traditionally, plans for missions have been developed over some period of time. Basic plans might have been developed over years of war gaming and can be adjusted in a matter of weeks or days to fit specific scenarios. As Helmuth von Moltke the Elder said, however, "no battle plan survives contact with the enemy" (Armour 2005). A hierarchical command structure consisting of the admiral in charge of the CSG, the commanding officers of the air wing and ships, and their subordinate structures allows for agile adjustments of plans. However, given the complexity of the domain and limited heuristics, decisions are often far from optimal.

Modern, net-centric technology allows for the sharing of sensor data so that all ships can have a common operational picture. Rather than helping the top-level decision makers, however, this wealth of information can often be overwhelming without adequate decision support tools.

Key Problems and Challenges for Distributed Planning

The CSG problem and other large-scale distributed systems such as logistics and parcel delivery exhibit some overlapping problems and challenges:

1. Dynamic environment: The environment changes (e.g., weather and enemy actions).
2. Dynamic goals: Mission priorities and objectives change.
3. Uncertainty: Actions have uncertain outcomes, as in the case of weapon systems malfunction, missed targets, or targets that employ counter-measures to avoid detection.
4. Scalability: Many independent entities must be controlled and coordinated. Ships and aircraft in the CSG are on the order of tens of entities. Each one of those entities has autonomous subsystems that may be interdependent with numerous other subsystems.

In the next section, we describe the Fortunate Artists Media Empire (FAME) domain. FAME is an artificial domain that allows us to focus and explore solutions to the core distributed planning issues while abstracting away some of the real-world complexities such as naval doctrine and its representation for the CSG problem.

The Fortunate Artists Media Empire (FAME) Problem Domain

FAME is a completely artificial distributed planning problem domain that incorporates the key problems and challenges described above: dynamic environments, dynamic goals, uncertainty of actions, and scalability. We first describe FAME, and then discuss how these key challenges are captured by this domain.

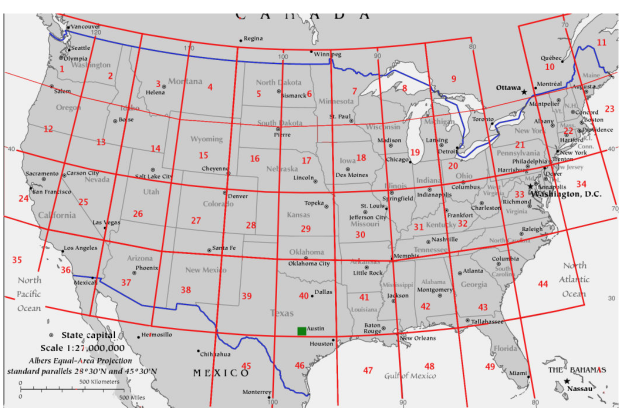

Figure 2: FAME map with a grid of 49 QA regional center (QARC) territories (UT Austin 2007).

The FAME Domain Description

We define the fictional Fortunate Artists Media Empire (FAME) quality assurance (QA) problem. FAME QA provides QA for thousands of radio and TV stations throughout the lower 48 states of the United States of America. It monitors broadcasts for Big Brother Organization (BBO) violations such as the use of obscene language. It also monitors advertising quotas. Its infrastructure forms a hierarchy that consists of a quality assurance headquarters (QAHQ), 49 QA regional centers (QARCs), and mobile, truck-mounted sensors used to monitor TV and radio broadcasts.

Quality Assurance Headquarters (QAHQ)

The Quality Assurance Headquarters (QAHQ) is located in Austin, Texas. It is responsible for developing a schedule of QA sampling and it receives sampling requests from advertisers or the BBO. Requests can also be generated as a result of monitored violations of advertising quotas or BBO regulations. There are two formats of sampling requests. One request is for a specific radio or TV station, a specific program, a specific timeslot and recurrence. This type of sampling request allows for broadcast delays if previous programs run over schedule. The other request format is for a specific station and time period. It monitors all programs during that time period. QAHQ determines which Quality Assurance Regional Centers (QARCs) can possibly cover the requests and makes assignments accordingly. For a centralized planning approach, QAHQ determines unique assignments. However, radio or TV stations that are located within sensor range of neighboring regions could be monitored from more than one region. In a distributed approach, a request could therefore be assigned to multiple QARCs and the QARCs would be responsible for deconflicting the assignments. Alternatively, all of the requests could be assigned to the entire group of mobile sensors in a territory, which would then need to coordinate in order to ensure the best coverage. Actual BBO or advertising violations are reported back to QAHQ via violation monitoring requests (VMRs).

Quality Assurance Regional Center (QARC)

Each Quality Assurance Regional Center (QARC) has a fixed geographical territory in which to employ its sensor assets. The bounding boxes for each QARC are the same size, but due to coastlines and borders some QARCs have a smaller landmass to cover than others. For example, Area 1 covers the US land-portion of the geographical area bounded by the latitude/longitude pairs 45N/130W and 50N/120W, while Area 2 is bounded by 45N/120W and 50N/110W. The bounding boxes of the 49 QARCs are shown in Figure 2. The bounding boxes are aligned to 5-degree increments in latitude and 10-degree increments in longitude. Each QARC has a fixed assigned number of mobile sensors. QARCs can also malfunction or have scheduled outages that would shut down all mobile sensors in a given region. In such cases, neighboring QARCs will attempt to cover high priority requests for neighboring radio or TV stations that could be monitored from their territory. QARCs also process monitored broadcasts received from mobile sensors for correctness of advertising quotas and BBO violations. For observed violations, QARCs generate high-priority VMRs, which are sent back to the QAHQ. For BBO violations, these VMRs are for a specific program for a specified length of recurrence, for example, a VMR to monitor every show of a specific disk jockey for one week. Advertising quota violations result in VMRs of the offending station for a specific length of time.

Mobile Sensors

Each QARC has a fixed number of assigned mobile, truck-mounted sensors. The mobile sensors must stay in the territory of their assigned QARC. Each sensor has a range that determines how close it must be to a radio or TV station in order to monitor its broadcast. Mobile sensors might be assigned to monitor stations in neighboring territories that are within their monitoring range. They have a maximum speed to transition to a new monitoring location but cannot monitor any broadcasts while moving.

Quality Assurance Methodology

There are scheduled QA samples and dynamic sampling requests of different priorities that must be covered by assigning collection sensor assets. If feasible, collection assets must monitor all VMRs. They should also attempt to monitor the highest-priority requests. The number of collection requests by far exceeds actual collection capabilities.

Communication and Broadcast Environment

QAHQ and the QARCs are connected via a high-speed wide-area network. All mobile sensors are uniquely named and can communicate directly with any QARC or other mobile sensor within their own region and neighboring regions. However, the communication infrastructure is primarily used for collection traffic that transports monitored broadcasts to the appropriate QARCs for QA processing and archival. Therefore, additional communication needed for command and control (C2) must be kept to a minimum. Due to the real-time nature of collection traffic, C2 traffic may be delayed or even dropped.

Our broadcast environment consists of radio and TV stations at known locations and published broadcast schedules. Each broadcast station has a probabilistic maximum broadcast range.

Quality Metrics

For this problem we propose two different quality metrics that determine the “goodness” of a solution:

1. Priority Weighting Metric (PWM): This metric provides different weights for each level of priority. Higher-priority requests have a much higher weight than lower-priority requests, but there are no penalties for sensor movement and the resulting losses of coverage.
2. PWM with Movement Penalty (PMP): This metric also uses priority weighting, but penalties are assigned for sensor movement and the resulting losses of coverage.

Experimental Parameters

The FAME domain provides a large number of adjustable parameters that can be used to map salient features from actual domains. The following is a non-exhaustive list of parameters that can be adjusted:

1. QAHQ Parameters: On the QAHQ level, we can change the number of sampling requests and the fixed QA sampling schedule. We can also control the geographic distribution of the sampling requests and their priorities.
2. QARC Parameters: For the regional level, we can adjust the number and the geography of the regions. We can model the frequency and distribution of QARC outages. The number of sensors could be the same for every region or proportional to the number of broadcast stations.
3. Mobile Sensor Parameters: Mobile sensors could have different monitoring capabilities. They could monitor either both, radio and TV broadcasts, or just one of those media. Sensors can also have different monitoring ranges. Other sensor parameters include probabilistic maximum speeds and probability of malfunction.
4. Communication Environment Parameters: We can model different probabilities of delayed or dropped traffic.
5. Broadcast Environment Parameters: We can adjust the number and geographic distribution of TV and radio stations. We can further model probabilistic broadcast distances and adherence to published broadcast schedules.

Mapping FAME Problems to Key Challenges

1. Dynamic environment: Quality Assurance Regional Center (QARC) outages, mobile sensor outages, and differing maximum speeds contribute to a highly dynamic environment.
2. Dynamic goals: With new monitoring requests of differing priorities, goals are constantly in flux. Monitored violations result in new high-priority requests.
3. Uncertainty: Mobile sensors have changing maximum speeds. Outages of mobile sensors and QARCs require replanning to cover highest-priority requests.
4. Scalability: The FAME domain allows us to experiment with large numbers of QARCs, and large numbers of mobile sensors per QARC. Different multiagent organizational paradigms like coalitions and hierarchies might allow us to implement more scalable solutions.

Summary and Future Work

We introduced the Fortunate Artists Media Empire (FAME) artificial problem domain, which allows us to examine distributed planning approaches in an abstract, dynamic environment that includes challenges such as dynamic goals, uncertainty, and scalability issues. In the FAME domain problem, we aim to allocate mobile, truck-mounted sensors to monitor broadcast stations according to prioritized monitoring requests. This domain provides a large number of experimentation parameters that can be mapped to salient features of actual domains. These features can then be explored against different distributed planning approaches. The domain task model—the monitoring requests—is a very flat task model consisting of only two subtasks, namely the potential movement of a mobile sensor followed by the actual monitoring subtask.

Our next step will be to model a realistic distribution of licensed broadcasting stations in the United States based on U.S. Federal Communications Commission data (U.S. Federal Communications Commission (FCC) 2006). Rather than implementing an ad-hoc model and simulation from scratch, we will use the JAMES II modeling and simulation framework (Himmelspach & Uhrmacher 2007), which supports modeling formalisms based on DEVS (Discrete EVent System specification) (Zeigler, Kim, & Praehofer 2000).

With sets of generated broadcast schedules and monitoring requests, we will then explore elements of distributed planning approaches such as the Distributed System for Interactive Planning and Execution (DSIPE) (desJardins & Wolverton 1999) and Generalized Partial Global Planning (GPGP) (Lesser *et al.* 2004), as well as distributed constraint optimization approaches such as Adopt (Modi *et al.* 2003) within that testbed. We also plan to explore multiagent organizational paradigms such as the use of dynamic coalitions (Horling & Lesser 2005), to increase scalability and survivability in pursuit of a hybrid distributed planning approach for the FAME domain.

For a class of distributed planning algorithms, the FAME domain could be used as a standardized environment and testbed that would allow the direct comparison of algorithms similar to RoboCup (Kitano *et al.* 1997) or the ACTIVE testbed designed for the DARPA Coordinators program (Emami *et al.* 2006).

References

Armour, P. G. 2005. To plan, two plans. *Communications of the ACM* 48(9):15–19.

desJardins, M., and Wolverton, M. 1999. Coordinating a distributed planning system. *AI Magazine* 20(4):45–53.

Emami, G.; Cheng, J.; Cornwell, D.; Feldhousen, M.; Long, C.; Malhotra, V.; Starnes, I.; Kerschberg, L.; Brodsky, A.; and Zhang, X. 2006. ACTIVE: Agile coordinator testbed integrated virtual environment. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '06)*, 1580–1587. New York, NY, USA: ACM Press.

Himmelspach, J., and Uhrmacher, A. 2007. Plug'n simulate. In *40th Annual Simulation Symposium (ANSS '07)*, 137–143. Los Alamitos, CA, USA: IEEE Computer Society.

Horling, B., and Lesser, V. 2005. A Survey of Multi-Agent Organizational Paradigms. *The Knowledge Engineering Review* 19(4):281–316.

Kitano, H.; Asada, M.; Kuniyoshi, Y.; Noda, I.; and Osawa, E. 1997. RoboCup: The robot world cup initiative. In Johnson, W. L., and Hayes-Roth, B., eds., *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, 340–347. New York: ACM Press.

Lesser, V.; Decker, K.; Wagner, T.; Carver, N.; Garvey, A.; Horling, B.; Neiman, D.; Podorozhny, R.; Prasad, N. M.; Raja, A.; Vincent, R.; Xuan, P.; and Zhang, X. Q. 2004. Evolution of the GPGP/TÆMS domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems* 9(1):87–143.

Modi, P. J.; Shen, W.-M.; Tambe, M.; and Yokoo, M. 2003. An asynchronous complete method for distributed constraint optimization. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems (AAMAS '03)*, 161–168. New York, NY, USA: ACM Press.

U.S. Federal Communications Commission (FCC). 2006. Broadcast station totals. http://www.fcc.gov/mb/audio/totals/bt061231.html. [Online; accessed on Sept. 11, 2007].

U.S. Navy Office of Information. 2004. The Carrier Strike Group. http://www.chinfo.navy.mil/. [Online; accessed 07 May 2007].

UT Austin. 2007. Based on map provided courtesy of the University of Texas Libraries. http://www.lib.utexas.edu/maps/united_states/usa_blank.jpg. [Online; accessed on 23 April, 2007].

Wikipedia. 2007. Carrier Strike Group — Wikipedia, The Free Encyclopedia. http://en.wikipedia.org/wiki/Carrier_battle_group. [Online; accessed 23 April 2007].

Zeigler, B. P.; Kim, T. G.; and Praehofer, H. 2000. *Theory of Modeling and Simulation*. Orlando, FL, USA: Academic Press, Inc.